

A Feasibility-Enforcing Primal-Decomposition SQP Algorithm for Optimal Vehicle Coordination

Mario Zanon, Robert Hult, Sébastien Gros and Paolo Falcone

Abstract—In this paper we consider the problem of coordinating autonomous vehicles approaching an intersection. We cast the problem in the distributed optimisation framework and propose an algorithm to solve it in real time. We extend previous work on the topic by testing two alternative algorithmic solutions in simulations. Moreover, we test our algorithm in experiments using real cars on a test track. The experimental results demonstrate the applicability and real-time feasibility of the algorithm.

I. INTRODUCTION

The problem of coordinating autonomous vehicles approaching an intersection has recently attracted increasing research interest [2], [4], [10]. The main motivation behind this research topic is to use automation in order to (a) reduce the amount of accidents, (b) reduce pollution and energy consumption and (c) increase the capacity of the infrastructure.

In order to achieve these three goals, it is necessary to introduce communication between the involved agents and the infrastructure, and to design a suitable algorithmic framework in order to compute a coordination policy. The coordination problem can be framed as a distributed mixed-integer optimal control problem. While this class of problems has a high complexity and is known to be NP-hard [3], by designing tailored optimisation algorithms and heuristics one can aim at computing approximated solutions in a rather short time. In this paper, however, we assume a prescribed crossing order and focus on the continuous part of the problem, i.e. the solution of the distributed optimal control problem. Future research will aim at extending the presented results in order to also optimise the crossing order.

The problem we address is not restricted to vehicles approaching an intersection, but can be framed more generally as a resource allocation problem subject to dynamic constraints. Such problems find applications in e.g. advanced manufacturing, traffic management, and logistics.

This paper is structured as follows. In Section II we introduce the formulation of the coordination problem as a distributed optimisation one. In Section III we present the distributed algorithm used to solve the problem. In Section IV we discuss the simulation and experimental results. We present the concluding remarks and outline for future research in Section V.

This work was supported by Copplar (project number 32226302), the Swedish Research Council (VR, grant number 2012-4038) and the European Commission Seventh Framework under the AdaptIVe (grant number 610428). The authors are with the Department of Signals and Systems, Chalmers University of Technology, Göteborg, Sweden. e-mail: {name.surname@chalmers.se}




Fig. 1. Illustration of the considered intersection scenario.

II. PROBLEM FORMULATION

We consider the problem of N_a vehicles approaching an intersection, as illustrated in Figure 1. For each vehicle i , we define the in- and out-times, i.e. the times at which the vehicle enters and exits the intersection, respectively as t_i^{in} and t_i^{out} . We lump them together in vectors $t_i := (t_i^{\text{in}}, t_i^{\text{out}})$ and we define the vector of all in- and out-times as $t := (t_1, \dots, t_{N_a})$.

The cost $V_i(t_i)$ associated with each individual vehicle is defined by an MPC problem specific to the vehicle. We define the states and controls of vehicle i at time k as $x_{i,k}$ and $u_{i,k}$ respectively. Moreover, we define vector $z_i := (z_{i,0}, \dots, z_{i,N-1}, x_{i,N})$ with $z_{i,k} := (x_{i,k}, u_{i,k})$ and the prediction horizon $N \in \mathbb{N}$. For simplicity, we assume linear dynamics and affine path constraints. Finally, we denote the set of all integers in a given interval as $\mathbb{I}_{[a,b]} = \{a, a+1, \dots, b\}$. The MPC problem of vehicle i is then given by

$$V_i(t_i) :=$$

$$\min_{z_i} J_i(z_i) \quad (1a)$$

$$\text{s.t. } x_{i,0} = \hat{x}_{i,0} \quad (1b)$$

$$x_{i,k+1} = A_i x_{i,k} + B_i u_{i,k}, \quad k \in \mathbb{I}_{[0,N-1]}, \quad (1c)$$

$$D_{i,k} x_{i,k} + E_{i,k} u_{i,k} \leq e_{i,k}, \quad k \in \mathbb{I}_{[0,N-1]}, \quad (1d)$$

$$p_i(t_i^{\text{in}}, z_i) - p_i^{\text{in}} = 0, \quad (1e)$$

$$p_i(t_i^{\text{out}}, z_i) - p_i^{\text{out}} = 0, \quad (1f)$$

where $J_i(z_i)$ is a vehicle-specific quadratic cost and constraints (1e) and (1f) make use of function p_i , describing the position of vehicle i on its path along road, in order to force

the vehicle to enter and exit the intersection at the prescribed times, see [9]. Because times t_i are fixed, Problem (1) is a Quadratic Program (QP). The set of feasible in- and out-times is then the domain of MPC Problem (1), i.e.

$$\mathcal{T}_i := \text{dom}(V_i(t_i)). \quad (2)$$

The coordination problem can then be formulated as

$$\min_t \sum_{i=1}^{N_a} V_i(t_i) \quad (3a)$$

$$\text{s.t. } t_i \in \mathcal{T}_i, \quad i \in \mathbb{I}_{[1, N_a]}, \quad (3b)$$

$$t_{i+1}^{\text{out}} \leq t_i^{\text{in}}, \quad i \in \mathbb{I}_{[1, N_a]}, \quad (3c)$$

such that the sum of each vehicle's individual cost is minimised, subject to having at most one vehicle in the intersection at any time (3c). Note that, because constraints (1e)-(1f) are generally nonlinear in t_i , Problem (3) is nonconvex.

We assume that the technical assumption [9, Assumption 1] holds and $p_i(t, z_i)$ is monotonically increasing in t and differentiable. Then, \mathcal{T}_i can be described by the in- out-times relative to the trajectories obtained by solving the linear programs (LPs):

$$(\max) \min_{z_i} x_N \text{ s.t. (1b) -- (1d), for bounds on } t_i^{\text{in}}, \quad (4a)$$

$$(\max) \min_{z_i} x_N \text{ s.t. (1b) -- (1e),} \quad (4b)$$

for t_i^{in} -dependent constraints on t_i^{out} , which we denote as $t_i^{\text{out}} \leq t_i^{\text{out}, \text{ub}}(t_i^{\text{in}})$, $t_i^{\text{out}} \geq t_i^{\text{out}, \text{lb}}(t_i^{\text{in}})$. For all details concerning this formulation we refer to [9].

III. NUMERICAL SOLUTION OF THE NLP

In this section, we present an adaptation of sequential quadratic programming (SQP) in a distributed setting for solving Problem (3). We remark that, by definition of $V_i(t_i)$ and \mathcal{T}_i , Problem (3) is non-smooth. A thorough discussion on the continuity and differentiability properties of V_i and \mathcal{T}_i has been presented in [9]. Note however that, provided that an interior-point QP solver is used in the local problems (1) and (4), one obtains a smooth approximation of the non-smooth problem. While this solves theoretical issues regarding the convergence of algorithms for smooth optimisation, slow convergence can in principle not be excluded, as the smoothly approximated problem becomes highly nonlinear at the points of non-smoothness of the original problem.

A. Sequential Quadratic Programming

For notational simplicity, we rewrite NLP (3) as

$$\min_t f(t) \quad \text{s.t. } h(t) \geq 0, \quad (5)$$

where $f(t) = \sum_{i=1}^{N_a} V_i(t_i)$ and we lump Constraints (3b)-(3c) in function h . We define the associated Lagrangian as $\mathcal{L}(t, \mu) := f(t) - \mu^\top h(t)$. Starting from an initial guess $v^{(0)} = (t^{(0)}, \mu^{(0)})$, SQP iteratively computes $v^{(j)}$ using

$$v^{(j+1)} = v^{(j)} + \alpha^{(j)} \Delta v^{(j)}, \quad (6)$$

with $\alpha^{(j)} \in (0, 1]$ and

$$\Delta v^{(j)} = (\Delta t^{(j)}, \tilde{\mu}^{(j)} - \mu^{(j)}), \quad (7)$$

obtained as the primal-dual solution $(\Delta t^{(j)}, \tilde{\mu}^{(j)})$ of the quadratic programming (QP) subproblem

$$\min_{\Delta t} \frac{1}{2} \Delta t^\top H^{(j)} \Delta t + \nabla f(t^{(j)})^\top \Delta t \quad (8a)$$

$$\text{s.t. } h(t^{(j)}) + \nabla h(t^{(j)})^\top \Delta t \geq 0. \quad (8b)$$

The iterations are stopped when the KKT residual r satisfies

$$r^{(j)} := \left\| \begin{array}{c} \nabla f(t^{(j)}) - \nabla h(t^{(j)}) \mu^{(j)} \\ \min(0, h(t^{(j)})) \end{array} \right\|_\infty \leq \epsilon. \quad (9)$$

Existing variants of SQP differ in the computations of the step size $\alpha^{(j)}$ and the so-called Hessian matrix approximation $H^{(j)}$. If exact Hessian is used, i.e. $H^{(j)} = \nabla_{tt}^2 \mathcal{L}(t^{(j)}, \mu^{(j)})$, the KKT conditions of each QP subproblem (8) coincide with a special form of linearisation of the KKT conditions of NLP (5) evaluated at $t^{(j)}$. For more details on SQP, we refer to e.g. [11].

The QP matrices $\nabla f(t^{(j)})$, $\nabla h(t^{(j)})$ and the Hessian of the Lagrangian $\nabla_{tt}^2 \mathcal{L}(t^{(j)}, \mu^{(j)})$ are called first and second-order sensitivities respectively, whose computation is presented in detail in [9]. Here we simply recall that the evaluations of f, h require the solution of 2 LPs and 1 QP, while the first and second order sensitivities are obtained at a marginal additional cost with respect to the function evaluations.

B. Hessian Regularisation

For nonconvex optimisation algorithms, it is in general required that the reduced Hessian is positive definite, in order to avoid solving indefinite QP subproblems. Therefore, a modification of the Lagrangian Hessian might be required, especially if the linearisation point is far from the optimum. For a thorough analysis we refer the interested reader to e.g. [11] and references therein.

In this paper, we apply the simple strategy of removing all directions of negative curvature present in the Hessian by adding the minimal regularisation needed to obtain a positive curvature in all directions. We do this by exploiting the block-diagonal structure of the Hessian to perform an eigenvalue decomposition of each 2-by-2 block and saturate all eigenvalues to a predefined minimum positive value. Given the small size, the eigenvalue decomposition can be made very efficient.

C. Globalisation and Merit Function

In order to guarantee global convergence of SQP algorithms, several approaches have been proposed in the literature, including *linesearch* and *trust region* methods, both of which rely on a so-called *merit function* to measure progress towards optimality and feasibility. Several merit functions have been proposed and we refer the interested reader to [11] and references therein for more details on the

topic. In this paper we make use of linesearch based on the so-called ℓ_1 merit function, defined as

$$M(t) = f(t) + \sigma \|h^-(t)\|_1, \quad (10)$$

where $h^-(t) := \min(h(t), 0)$ and σ is a parameter which must be chosen such that $\sigma > \|\mu\|_\infty$.

Note that one can compute the first-order term of the Taylor expansion of M in the SQP direction y as [11]:

$$DM(x)[y]y = \nabla f(x)y - \sigma \|h^-(x)\|_1, \quad (11)$$

where we define the directional derivative of M in direction y evaluated at x as $DM(x)[y]$.

Once the QP solution $\Delta v^{(j)}$ is obtained, linesearch globalisation techniques choose the step size so as to enforce a decrease in the merit function based on the Armijo condition

$$M(t + \alpha\Delta t) \leq M(t) + \gamma DM(t)[\Delta t]\alpha\Delta t, \quad (12)$$

where $\gamma \in (0, 0.5]$ is a fixed constant.

Because solving (12) exactly is in general computationally expensive, a popular technique called *backtracking* consists in starting with $\alpha = 1$ and iteratively reducing it using $\alpha \leftarrow \beta\alpha$ with $\beta \in (0, 1)$ until condition (12) is satisfied.

D. Local Feasibility Issues

NLP solvers only guarantee feasibility of the constraints at convergence and not throughout the iterations. Because in the case of Problem (3), the cost terms $V_i(t_i)$ are not defined for infeasible times $t_i \notin \mathcal{T}_i$, a remedy needs to be devised.

We compare two methods based on: (a) introducing soft constraints using slack variables and penalising the ℓ_1 -norm of the constraint violation with a sufficiently high weight and (b) using a feasibility-restoring projection of each iterate, as proposed in [13]. In the following we will call approach (a) *relaxation* and approach (b) *projection*. In [13] the use of projection is preferred to the use of relaxation. However, this choice is not motivated formally and no comparison between the two approaches is made.

While appealing for its simplicity, approach (a) suffers from one main drawback: the local quadratic approximation evaluated at an infeasible point is dominated by the cost associated with the constraint violation and loses its validity once the problem becomes feasible. This approach can be implemented by replacing constraints (1f) in each agent's problem by

$$p_i(t_i^{\text{out}}, z_i) - p_i^{\text{out}} + s_1 - s_2 = 0, \quad (13a)$$

$$s_1 \geq 0, \quad s_2 \geq 0. \quad (13b)$$

Proposition 1: Reformulate constraints (1f) as (13) and to add the term $\rho s_1 + \rho s_2$ to the cost of each local Problem (1). Choose $\rho > \lambda_{\max}$, i.e. the largest Lagrange multiplier associated with any constraint of the local Problem (1) for all feasible in-out times t_i . Then, the SQP algorithm with linesearch will converge.

Proof: The condition $\rho > \lambda_{\max}$ is necessary in order to make sure that the relaxed problem yields the solution to the original problem whenever the latter has a

feasible solution [6, Theorem 14.3.1]. The relaxation of the constraints ensures that the local problems are feasible for any in- out-time pair and linesearch ensures a decrease in the merit function. Therefore, for all feasible t , the relaxed formulation coincides with the original one, while for all infeasible t , linesearch ensures that progress is made towards both feasibility and optimality. ■

In order to efficiently implement the projection approach (b), we observe that (i) the constraints (3b) are simple bounds on the in times t_i^{in} which are always satisfied, and (ii) the constraints on the out times are given by $t_i^{\text{out}} \geq t_i^{\text{olb}}(t_i^{\text{in}})$ and $t_i^{\text{out}} \leq t_i^{\text{out,ub}}(t_i^{\text{in}})$. Therefore, the in times t_i^{in} are always feasible and all out times t_i^{out} which become infeasible at iterate k can be projected back onto the set of feasible times by using

$$\mathcal{P}_i(t_i^{\text{out}(j)}) := \left[t_i^{\text{out,ub}}(t_i^{\text{in}(j)}) \right]_{t_i^{\text{olb}}(t_i^{\text{in}(j)})}^{t_i^{\text{out}(j)}},$$

with $[a]_b^c := \max(\min(a, c), b)$. Moreover, such a projection can be done easily without extra computations, as the evaluation of $t_i^{\text{out,ub}}(t_i^{\text{in}(j)})$, $t_i^{\text{olb}}(t_i^{\text{in}(j)})$ is needed for each iterate $j + 1$ of the SQP algorithm. An illustration of the projection procedure is given in [13]. The main drawback of this approach is the partial loss of parallelisability at the agent level, i.e. the two LPs for the computation of the linearisation of constraint (3b) must be solved before QP (1).

While the projection is simple to perform, the linesearch procedure has to be modified in order to guarantee that the step is a descent direction, provided that α is chosen small enough. We therefore define a modification of the Armijo condition (12) using the projection as follows

$$M(\mathcal{P}(t + \alpha\Delta t)) \leq M(t) + \gamma DM(t)[\Delta t]\alpha\Delta t, \quad (14)$$

with Δt computed by (7). As opposed to the standard linesearch approaches, in (14) we do not backtrack from the projected Newton step. Instead, we backtrack from the non-projected Newton step to obtain the step candidate and project each step candidate.

Proposition 2: The Armijo condition with projection (14) is guaranteed to hold, provided that α is chosen small enough.

Proof: The proof is provided in [13]. ■

IV. RESULTS

In order to test the algorithm, we have considered vehicles defined by the linear system

$$\dot{x} = A_c x + B_c u, \quad A_c = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B_c = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

where $u \in [u^{\text{lb}}, u^{\text{ub}}]$ m/s² is the (scalar) control and $x = (p, v)$ is the state vector, which includes position p and velocity $v \geq 0$. We discretise the system using zero-order hold to obtain

$$x_{k+1} = A x_k + B u_k, \quad A = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.5 T_s^2 \\ T_s \end{bmatrix},$$

TABLE I
SIMULATION PARAMETERS (SI UNITS)

Vehicle #	1	2	3	4	5	6
Q_i	1	1	10	10	1	1
R_i	1	1	1	1	1	1
v_i^d	80	80	65	70	70	60
\hat{p}_0	-55	-60	-55	-70	-70	-60
\hat{v}_0	80	80	65	70	70	60

TABLE II
SCENARIO CROSSING ORDER

# scenario	1	2	3	4	5	6	7
Order	1	1	2	1	1	2	2
	2	2	1	2	2	1	1
	6	4	3	5	3	3	3
	3	3	4	4	4	4	6
	4	5	5	3	6	6	4
	5	6	6	6	5	5	5

where T_s is the sampling time. Each vehicle i has the cost

$$J_i(z_i) = \sum_{k=0}^N Q_i(v_{i,k} - v_{i,k}^d)^2 + R_i u_{i,k}^2, \quad (15)$$

where Q_i and R_i are vehicle-specific weights and $v_{i,k}^d$ is the desired speed of the vehicle at time instant k .

A. The MPC Formulation

In the previous sections, we have presented the algorithm to compute the in-out times. While in principle one could formulate the MPC problem of each vehicle based on Problem (1), we preferred to adopt a slightly modified version, where constraints (1e)-(1f) have been relaxed as $p_i(t_i^{\text{in}}, z_i) \leq p_i^{\text{in}}, p_i(t_i^{\text{out}}, z_i) \geq p_i^{\text{out}}$. Note that this relaxation still enforces the safety-critical collision avoidance constraint. Moreover, in order to retain feasibility even in the presence of sensor noise, we formulated constraints (1e)-(1f) as soft constraints with exact penalty [5], [12].

B. Simulations

All simulation parameters are specified in Tables I and II, where the units have been omitted for ease of reading. The intersection is defined by $p_i^{\text{in}} = 0$ m, and $p_i^{\text{out}} = 8$ m and the control bounds are given by $-u^{\text{lb}} = u^{\text{ub}} = 2$ m/s² for each vehicle i .

We compared the two feasibility-enforcing approaches by means of numerical simulations. We have solved 5 scenarios using both the projection and the slack-based feasibility-enforcing techniques. In all simulations we used the parameters $\gamma = 0.01$, $\beta = 0.5$, $\epsilon = 10^{-2}$.

We stress again that evaluating functions f and h constitutes the major computational cost (solving QP (1) and the two LPs (4)), while the sensitivity information ∇f , ∇h , $\nabla_{tt}^2 \mathcal{L}$ is much cheaper to compute. Therefore, every linesearch iteration has about the same computational cost as one SQP iteration. On the communication side, every linesearch iteration is a bit cheaper than an SQP iteration, as much less information needs to be communicated.

TABLE III
SQP AND BACKTRACKING LINESEARCH ITERATIONS

# scenario	1	2	3	4	5	6	7
Projection	n_{SQP}	6	10	5	4	7	8
	n_{ls}	6	10	5	4	7	8
Relaxation	n_{SQP}	5	7	9	16	9	8
	n_{ls}	7	10	15	49	8	16

The results are displayed in Table III. It can be seen that in some scenarios relaxation and projection yield comparable results in terms of SQP and linesearch iterates n_{SQP} and n_{ls} respectively. However, in some scenarios the relaxation approach performs much worse than the projection approach. We moreover remark that the relaxation approach never required Hessian regularisations, while the projection approach added some regularisation once in Scenarios 1 and 4 and twice in Scenario 7.

C. Experiments

In this subsection, we present the results of experiments performed at the AstaZero proving ground next to Gothenburg, Sweden. A video summarising the experimental tests is available at [8]. In these experiments, the coordination algorithm was validated in closed loop on a real intersection using three automated Volvo cars (two Volvo S60 sedans and one Volvo XC90 SUV) and vehicle-to-vehicle (V2V) communication from RENDITS [1]. Each vehicle was further equipped with Real-Time-Kinematic (RTK) capable GPS receivers, inertial sensors, and a computer network consisting of a MicroAutobox 2 real-time computer and a laptop, on which the experiment software was executed. The SQP algorithm was applied in a distributed fashion using the V2V communication, where the SQP subproblems (8) computing the primal-dual updates were solved at a central computational node. The evaluation of the cost function, constraints and the corresponding sensitivities were done on-board the vehicles by solving the QPs and LPs using HPMPC [7]. Due to its higher reliability, the projection approach was chosen for the experimental implementation.

In each experiment the vehicles were controlled from stand-still to a configuration from which a three way collision would occur if the speeds of all vehicles were held constant. The coordination was set to start at a predefined time. Short before, the NLP was solved using the proposed distributed algorithm by relying on wireless communication between the vehicles and a central node.

Over the experimental campaign, we successfully completed over 100 runs with 3 vehicles and over 50 runs with 2 vehicles, and explored various parameter settings. However, for the sake of brevity, we present the results of only one run in this paper. In this run, the vehicles were all controlled to be at $\hat{p}_0 = -200$ m, $\hat{v}_0 = 50$ km/h at time $t_c = 31.3$ s from the beginning of the run, which was also used as the starting time for the coordination. Furthermore, for all vehicles, the control objective was to track the desired velocity $v^d = 50$ km/h, using the weights $Q_i = 1 \frac{\text{s}^2}{\text{m}^2}$ and $R_i = 10 \frac{\text{s}^4}{\text{m}^2}$ and the horizon length $N = 200$. The control bounds were set to




Fig. 2. SQP convergence and optimal times: first car in blue, second in red and third in black. Top graph: in-out times for the three vehicles. Bottom graph: norm of the KKT residual r (black) and step Δt (red).

$u^{\text{lb}} = -3 \text{ m/s}^2$, $u^{\text{ub}} = 1.6 \text{ m/s}^2$. Finally, we used the SQP parameters $\gamma = 0.01$, $\beta = 0.5$, $\epsilon = 10^{-2}$.

The results are displayed in Figure 2, where one can see the evolution in time of the KKT residual r , the step size Δt and the solution at the given time. In this particular run, the solver always takes full steps without any projection, which entails that feasibility is obtained after one step while the successive steps improve optimality. All times are given with respect to the beginning of the scenario, which includes the startup phase to bring the vehicles to the assigned initial configuration from standstill. While the SQP algorithm is initialized at $t = 24.0$ s, the sensitivities for the first iterate become available only at $t = 24.5$ s. To compute the SQP iterates, the central node needs to wait 5, 10, 5 and 5 sampling instants respectively. These comparatively long iteration times are due to implementation-specific details and could be reduced significantly with a better implementation. In particular, we emphasize that the total computation time, i.e. solving the central SQP subproblems and the vehicle level QPs and LPs, only constitutes about 3.4% of the total computation time. Moreover, the total time per iterate varies over the iterations. This is due to the unreliability of the wireless communication channel, e.g. some packets are dropped. Therefore, variations would be present even with a more efficient implementation, though with a much smaller magnitude.

The closed-loop trajectories are displayed in Figure 3. As one can expect, at the solution $t_1^{\text{out}} = t_2^{\text{in}}$ and $t_2^{\text{out}} = t_3^{\text{in}}$. Instead of proceeding at constant velocity, the velocities are adjusted and the vehicles deviate from their reference in order to meet the crossing times. Note that the second car does not keep a constant velocity but accelerates as it approaches the intersection. This reduces the duration of its permanence inside the intersection, thus allowing the other cars to deviate less from their references and, therefore, to reduce the overall cost.

Without coordination, all vehicles would proceed from




Fig. 3. Position trajectories (top graph, a zoom is given in Figure 4): first car in blue, second in red and third in black. The intersection start and stop is marked by two black lines, the coordination starts when the three vehicles are at $\hat{p}_0 = -200$ m and the vertical lines mark the in- out-times for each vehicle. Deviation of the position with respect to the constant speed scenario (middle graph). Velocity trajectory (bottom graph): reference velocity v^d in dashed lines and measured velocities in continuous line.

the coordination starting point at a constant velocity $v = 50$ km/h. The difference in position between the coordinated and uncoordinated solutions, here denoted p_{50} , is also displayed Figure 3. Because we do not provide any reference for the position, after crossing the intersection, the first car is ahead of the position it would have in the absence of the other two cars, whereas the last vehicle lags behind.

The positions shown in Figure 3 and Figure 4 were computed based on readings from GPS receivers, and the quality of the signal differs between three cars. In particular, the GPS unit for the third car (red line in the plots) systematically provided unreliable measurements with a frequency of 1 Hz. Due to such measurement errors and the mismatch between the prediction model and the physical vehicles, the constraints enforcing the in and out times are always violated by the closed loop system. We define the time violations as δt_i^{in} , δt_i^{out} , where a positive sign means that the actual time was larger than planned, such that $\delta t_i^{\text{in}} \geq 0$ corresponds to a safe configuration (the vehicle entered late), while $\delta t_i^{\text{out}} \geq 0$ corresponds to a dangerous configuration (the vehicle left late). For the position we kept the convention of the constraint definition, such that $p_i(t_i^{\text{in}}, z_i) - p_i^{\text{in}} \leq 0$ is safe, while $p_i(t_i^{\text{out}}, z_i) - p_i^{\text{out}} \leq 0$ is dangerous. The estimated violation of the in-out times for the experiment run is given in Table IV. It should be noted that the values given in Table IV are taken from the position signal displayed in Figure 4. For vehicle 2 and 3 this signal is rather noisy and not very reliable, hence the large violations. We remark that the observed constraint violations can be easily accounted for in the proposed setup by constraint tightening, i.e. by enlarging the intersection definition. Because the constraint has the dimension of a distance, the tightening procedure can be directly related to positioning errors.

The analysis of the obtained experimental results suggests




Fig. 4. Zoom of the position trajectories from Figure 3 in the intersection area. The horizontal dashed lines mark the beginning and the end of the intersection for each vehicle; the red line is superposed to the black one, as the two corresponding cars are identical. The vertical lines mark the in and out times sent to each car by the central node.

that the proposed decomposition of the control scheme has the following desirable properties: (a) it does not need to rely on efficient and reliable communication; (b) even if not implemented in the most efficient way, the algorithm converges fast enough to allow for a real-time implementation; (c) after the solution has been computed, communication is not required anymore and the vehicles can be safely controlled in a decoupled fashion; (d) sporadic wrong sensor readings are natively handled by the algorithm.

After a thorough experimental campaign in a safe configuration, i.e. driving the vehicles on parallel lanes, we ran some experiments on a real crossing scenario. For safety reasons, we added 5 m to the intersection length. An aerial picture is displayed in Figure 5, where we highlighted the intersection by a red box. In the real crossing experiment, we obtained results comparable to those obtained in the parallel lane configuration.

V. CONCLUSIONS AND FUTURE RESEARCH

In this paper, we have proposed a distributed control scheme for optimal vehicle coordination at intersections. We have developed an ad-hoc optimisation algorithm based on primal decomposition and a feasibility-enforcing projection. We have compared this approach versus a the relaxation feasibility-enforcing approach in simulations which have indicated a lower reliability of the latter.

We have tested our control approach in experiments using real cars on a test track. Though our implementation of the algorithm on the real-time platform was simple and far from optimal, the experimental results have demonstrated the applicability and real-time feasibility of our approach. Moreover, the sporadic faulty sensor readings and the unreliable communication have both demonstrated some degree of robustness of our control approach.

Future work will consider improved approaches that (a) update the in and out times in a closed-loop fashion, (b) solve




Fig. 5. Aerial picture of the crossing experiment.

TABLE IV
IN-OUT TIME AND POSITION ERRORS (SI UNITS)

Car #	δt_i^{in}	δt_i^{out}	$p_i(t_i^{\text{in}}, z_i) - p_i^{\text{in}}$	$p_i(t_i^{\text{out}}, z_i) - p_i^{\text{out}}$
1	0.0287	0.0078	-0.4363	-0.1213
2	0.0077	-0.0156	-0.2461	0.2914
3	0.0480	0.0257	-0.6610	-0.5217

the ordering problem, i.e. the full mixed-integer problem.

REFERENCES

- [1] RENDITS. <http://www.rendits.com/>. Accessed: 2016-11-05.
- [2] S.R. Azimi, G. Bhatia, R. Rajkumar, and P. Mudalige. Vehicular networks for collision avoidance at intersections. *SAE International Journal of Passenger Cars - Mechanical Systems*, 4(1):406–416, 2011.
- [3] A. Colombo and D. Del Vecchio. Efficient algorithms for collision avoidance at intersections. In *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control*, pages 145–154, New York, NY, USA, 2012.
- [4] G. R. de Campos, P. Falcone, and J. Sjöberg. Autonomous cooperative driving: A velocity-based negotiation approach for intersection crossing. In *Intelligent Transportation Systems - (ITSC), 2013 16th International IEEE Conference on*, pages 1456–1461, 2013.
- [5] N.M.C. de Oliveira and L.T. Biegler. Constraint Handling and Stability Properties of Model-Predictive Control. *AIChE Journal*, 40(7):1138–1155, 1994.
- [6] R. Fletcher. *Practical Methods of Optimization*. Wiley, Chichester, 2nd edition, 1987.
- [7] G. Frison, H.B. Sørensen, B. Dammann, and J.B. Jorgensen. High-performance small-scale solvers for linear model predictive control. In *Proc. 2014 European Control Conference (ECC)*, pages 128–133, June 2014.
- [8] R. Hult, M. Zanon, S. Gros, and P. Falcone. Optimal coordination of three cars approaching an intersection. <https://youtu.be/nYSXvnaNRK4>, 2016. Accessed: 2017-03-03.
- [9] R. Hult, M. Zanon, S. Gros, and P. Falcone. Primal Decomposition of the Optimal Coordination of Vehicles at Traffic Intersections. In *Proceedings of the 55th IEEE Conference on Decision and Control*, 2016. (accepted).
- [10] J. Lee and B. Park. Development and evaluation of a cooperative vehicle intersection control algorithm under the connected vehicles environment. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):81–90, 2012.
- [11] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2 edition, 2006.
- [12] P.O.M. Scolaert and J.B. Rawlings. Feasibility Issues in Linear Model Predictive Control. *AIChE Journal*, 45(8):1649–1659, 1999.
- [13] M. Zanon, S. Gros, P. Falcone, and H. Wymeersch. An Asynchronous Algorithm for Optimal Vehicle Coordination at Traffic Intersections. In *20th IFAC World Congress*, 2017. (accepted), preprint available at: <https://mariozanon.wordpress.com/asynchronous/>.